

# How do we remember the past in randomised strategies?

Julien Cristau

LIAFA  
CNRS & Université Denis Diderot - Paris 7  
Paris, France  
jcristau@liafa.jussieu.fr

Claire David

LFCS  
University of Edinburgh  
Edinburgh, Scotland  
cdavid@inf.ed.ac.uk

Florian Horn

LIAFA  
CNRS & Université Denis  
Diderot - Paris 7  
Paris, France  
horn@liafa.jussieu.fr

IST Austria  
Institute of Science and  
Technology Austria  
Klosterneuburg, Austria

Graph games of infinite length are a natural model for open reactive processes: one player represents the controller, trying to ensure a given specification, and the other represents a hostile environment. The evolution of the system depends on the decisions of both players, supplemented by chance.

In this work, we focus on the notion of *randomised* strategy. More specifically, we show that three natural definitions may lead to very different results: in the most general cases, an almost-surely winning situation may become almost-surely losing if the player is only allowed to use a weaker notion of strategy. In more reasonable settings, translations exist, but they require infinite memory, even in simple cases. Finally, some traditional problems become undecidable for the strongest type of strategies.

*“You can’t have a strategy against telepaths: you have to act randomly. You have to not know what you’re going to do next. You have to shut your eyes and run blindly. The problem is: how can you randomise your strategy, yet move purposefully towards your goal?”*

Solar Lottery  
Philip K. Dick

## 1 Introduction

Since their introduction to verification in the late eighties, graph games have emerged as the model of choice for problems about *open systems*, where a controller (Eve) must interact with an *a priori* hostile environment (Adam) [PR89]. In such games, an *arena* —i.e. a graph— models the system and its evolution: at the beginning, a token is laid on one of the vertices, and its moves are determined by the actions of the players, supplemented by chance. The infinite sequence of vertices that ensues constitutes a *play* of the game, whose winner is defined by some predetermined specification, often given as a regular condition on infinite words [MP92].

This model has been declined in a multiplicity of variants, in terms of both arenas and objectives. However, the questions are nearly always the same: Is there a winning strategy? For which player? How complex is it, in terms of memory and randomisation? Memory is the quantity of information that one is allowed to remember from the past: in general, the whole history is available, but it is often enough to

remember a finite quantity of information. In addition, a strategy is *pure* if it proposes only one possible action after any given sequence of observations. The notion of *randomised* strategy, and its relation to memory, is the subject of this paper.

In verification, “randomised strategy” usually refers to a function from the history to a probability distribution over the actions. In other domains, such strategies are called “*behavioural* strategies”, as opposed to two other models of randomised strategies: a *mixed* strategy is a measure over pure strategies, and a *general* strategy is a measure over behavioural strategies. These models are also relevant in computer science. Indeed, the IPv6 “Stateless Address Autoconfiguration” protocol, which only uses randomisation at the beginning to generate a new I.P. address [TNJ07], can be accurately described as a mixed strategy. Likewise, the secure shell protocol (ssh) is a general strategy, since a new session key has to be randomly<sup>1</sup> generated every hour or gigabyte [YL06].

In this paper, we propose definitions for mixed and general strategies, with or without memory, in the framework of graph games for verification. We expose several situations in which their analyses differ significantly from the behavioural model. In the most general case, the same game can be almost-surely losing or almost-surely winning depending on the type of strategies we consider. In other situations, we conjecture that the values are the same, but we show that memory needs vary (from two to infinity). Altogether, we hope to ask more questions than we give answers: our main objective is to describe these three models for randomised strategies in graph games and to point out that many problems which are solved for behavioural strategies are still open in the mixed and general cases.

The paper is organised as follows. In Section 2, we recall the classical notions about graph games in verification, in a very general framework which subsumes a large part of the literature. Section 3 presents our definitions for behavioural, mixed, and general strategies in graph games, and stresses the fundamental differences between the three notions. Section 4 focuses on memory-related issues: it exhibits variations in the elementary cases of concurrent safety games and simple Muller games. In Section 5, we sum up our observations and results, and propose some open problems.

## 2 Definitions

### Notation

For a finite or countable set  $\mathcal{S}$ , we denote by  $\mathcal{D}(\mathcal{S})$  the set of probability distributions over  $\mathcal{S}$ , *i.e.* the set of functions from  $\mathcal{S}$  to positive real numbers that sums up to one.

### Arenas and plays

An *arena*  $\mathcal{A}$  is a tuple  $(\mathcal{V}, \mathcal{X}, \mathcal{Y}, \delta, \Gamma, \gamma, \Phi, \varphi, \Psi, \psi)$  where  $\mathcal{V}$  is the set of vertices in the graph,  $\mathcal{X}$  is the set of actions of Eve,  $\mathcal{Y}$  is the set of actions of Adam,  $\delta : \mathcal{V} \times \mathcal{X} \times \mathcal{Y} \rightarrow \mathcal{D}(\mathcal{V})$  is the transition function,  $\Gamma$  is the set of colours,  $\gamma : \mathcal{V} \rightarrow \Gamma$  is the colouring function,  $\Phi$  is the set of signals of Eve,  $\varphi : \mathcal{V} \cup \mathcal{X} \rightarrow \Phi$  is her observation function,  $\Psi$  is the set of signals of Adam, and  $\psi : \mathcal{V} \cup \mathcal{Y} \rightarrow \Psi$  is his observation function. Many results about graph games for verification consider only restricted arenas, such as:

**Synchronous:** an arena is *synchronous* if  $\varphi$  and  $\psi$  are total.

**Observable actions:** a synchronous arena has *observable actions* if the restrictions of  $\varphi$  to  $\mathcal{X}$  and  $\psi$  to  $\mathcal{Y}$  are one-to-one.

---

<sup>1</sup>except on Debian [BB08].

**Perfect information:** a synchronous arena has *perfect information* (or is *concurrent*) if the restrictions of  $\varphi$  and  $\psi$  to  $\mathcal{V}$  are one-to-one.

**Simple:** a concurrent arena is *simple* (or *turn-based*) if for each vertex  $q$ ,  $\delta(q, x, y)$  depends either on  $x$  or on  $y$ , but not both.

A *play* on the arena  $\mathcal{A}$  is a (possibly infinite) sequence  $\pi = \pi_0 \pi_1 \dots$  of states such that  $\forall i < |\pi| - 1, \exists x \in \mathcal{X}, y \in \mathcal{Y}, \delta(\pi_i, x, y)(\pi_{i+1}) > 0$ . The set of plays is usually denoted  $\Omega$ , and the set of plays starting with the vertex  $q$  by  $\Omega_q$ .

### Pure strategies and measures

A *pure strategy*  $\sigma$  for Eve (resp.  $\tau$  for Adam) on the arena  $\mathcal{A}$  associates an action to each finite sequence of observations:  $\sigma : \Phi^* \rightarrow \mathcal{X}$  (resp.  $\tau : \Psi^* \rightarrow \mathcal{Y}$ ). A play  $\pi$  is *consistent with a strategy*  $\sigma$  for Eve (resp.  $\tau$  for Adam) if and only if at each step  $i$ , there is an action  $y \in \mathcal{Y}$  for Adam (resp.  $x \in \mathcal{X}$ ) such that  $\delta(\pi_i, \sigma(\varphi(\pi_{1..i})), y)(\pi_{i+1}) > 0$  (resp.  $\delta(\pi_i, x, \tau(\psi(\pi_{1..i}))) (\pi_{i+1}) > 0$ ). Notice that, in the case of an asynchronous arena, actions can only change with new observations: otherwise, the same argument leads to the same result over and over. The set of plays consistent with  $\sigma$  (resp.  $\tau$ ;  $\sigma$  and  $\tau$ ) is denoted by  $\Omega^\sigma$  (resp.  $\Omega^\tau$ ;  $\Omega^{\sigma, \tau}$ ). Once an initial vertex  $q$  and two strategies  $\sigma$  and  $\tau$  have been fixed,  $\Omega_q^{\sigma, \tau}$  can naturally be made into a measurable space  $(\Omega_q^{\sigma, \tau}, \mathcal{O})$ , where  $\mathcal{O}$  is the  $\sigma$ -field generated by the cones  $\{\mathcal{O}_w \mid w \in \mathcal{V}^*\}$ :  $\pi \in \mathcal{O}_w$  if and only if  $w$  is a prefix of  $\pi$ . The probability measure  $\mathbb{P}_q^{\sigma, \tau}$  is recursively defined by  $\mathbb{P}_q^{\sigma, \tau}(\mathcal{O}_q) = 1$  and:

$$\forall w \in \mathcal{V}^*, (r, s) \in \mathcal{V}^2, \mathbb{P}_q^{\sigma, \tau}(\mathcal{O}_{wrs}) = \mathbb{P}_q^{\sigma, \tau}(\mathcal{O}_{wr}) \cdot \delta(r, \sigma(\varphi(wr)), \tau(\psi(wr)))(s).$$

Carathéodory's extension theorem allows us to extend  $\mathbb{P}_q^{\sigma, \tau}$  to the Borel sets of  $(\Omega_q^{\sigma, \tau}, \mathcal{O})$  [Wil91].

### Winning conditions and values

A *winning condition*  $\mathcal{W}$  on a set of colours  $\Gamma$  is a Borel subset of  $\Gamma^\infty$ . A play  $\pi$  in an arena  $\mathcal{A}$  on  $\Gamma$  is *winning for Eve* in the game  $(\mathcal{A}, \Gamma)$  if  $\gamma(\pi) \in \mathcal{W}$ , and *winning for Adam* otherwise. In a game  $\mathcal{G} = (\mathcal{A}, \mathcal{W})$ , the *pure value of a state*  $q$  under the strategies  $\sigma$  and  $\tau$ , denoted  $v_{\sigma, \tau}(q)$ , is the measure of  $\mathcal{W}$  under  $\mathbb{P}_q^{\sigma, \tau}$ . The *value for Eve of a state*  $q$  is the supremum of the values that she can ensure from  $q$  against any strategy of Adam. Symmetrically, the *value for Adam of a state*  $q$  is the infimum of the values he can defend against any strategy of Eve. In simple stochastic games, these two values coincide and are usually called the *value of*  $q$  [Mar98, MS98]<sup>2</sup>.

$$\mathbf{v}(q) = \sup_{\sigma} v_{\sigma}(q) = \inf_{\tau} v_{\tau}(q) .$$

### Winning criteria

Following de Alfaro and Henzinger [dAH00], we consider several notions of *winning strategies* and *winning regions*, depending on the chances Eve has to win. In decreasing order of difficulty, and from an initial vertex  $q$ , a strategy  $\sigma$  for Eve:

- is *sure* if any play consistent with  $\sigma$  is winning for Eve;

---

<sup>2</sup>As a matter of fact, these papers shows the quantitative determinacy, in behavioural strategies, of Borel games on concurrent arenas. An inspection of the proof yields the same result for pure strategies in the case of simple arenas.

- is *almost-sure* if for any strategy  $\tau$  for Adam,  $v_{\sigma,\tau}(q) = 1$ ;
- *ensures*  $\varepsilon$  if for any strategy  $\tau$  for Adam,  $v_{\sigma,\tau}(q) \geq \varepsilon$ ;
- is *positive* if for any strategy  $\tau$  for Adam,  $v_{\sigma,\tau}(q) > 0$ ;
- is *heroic* if for any strategy  $\tau$ , there is a play consistent with  $\sigma$  and  $\tau$  which is winning for Eve.

The *sure region* (resp. *almost-sure region of Eve*, *positive region*, *heroic region*) of Eve is the set of vertices from which she has a sure (resp. almost-sure, positive, heroic) strategy. Furthermore, the *bounded region* is the set of vertices from which Eve has a strategy ensuring a positive  $\varepsilon$  and the *limit-one region* is the set of vertices from which Eve has a strategy ensuring  $\varepsilon$  for any  $\varepsilon < 1$ . The same concepts are defined accordingly for Adam, except that we say that a strategy  $\tau$  for Adam *defends*  $\varepsilon$  if it guarantees that, for any strategy  $\sigma$  for Eve,  $v_{\sigma,\tau}(q) \leq \varepsilon$ .

### 3 Behavioural, mixed, and general strategies

As soon as we deal with concurrent arenas, we cannot rely only on pure strategies to make meaningful analyses. In the classical game of “Janken”, any pure strategy is surely beaten by the appropriate counter-strategy (paper against rock, scissors against paper, and rock against paper), but a strategy which plays each action with probability  $\frac{1}{3}$  eventually wins with probability  $\frac{1}{2}$ . The main point of this paper is that there are several possible definitions for the notion of “randomized strategy”.

- A *behavioural strategy* returns at each step a distribution over the actions:  $\Phi^* \rightarrow \mathcal{D}(\mathcal{X})$ ;
- a *mixed strategy* is a measure over pure strategies:  $\mathcal{D}(\Phi^* \rightarrow \mathcal{X})$ ;
- a *general strategy* is a measure over behavioural strategies:  $\mathcal{D}(\Phi^* \rightarrow \mathcal{D}(\mathcal{X}))$ .

As we show in this paper the expressive powers of these models are quite different. Intuitively, a behavioural strategy does not know in advance what it will play next, so its actions can change when its decisions do not (even when there are no observations). Mixed strategies use randomization to get hidden information at the beginning of the play, which can later be used to correlate undistinguishable actions, e.g. playing *aa* or *bb* with probability  $\frac{1}{2}$ . General strategies subsume both, so they can, in particular, generate hidden information on the fly. These distinctions have mostly been overlooked in verification (apart from a few remarks, e.g. [dAHK98, DHR08]). One reason is that the games we consider are usually synchronous, with observable actions. On synchronous arenas, mixed strategies can simulate behavioural strategies: as each action can be uniquely identified beforehand by its position in the play, it is possible to define a measure which somehow makes all the random draws at the beginning of the play. If furthermore, the actions are observable, Kuhn’s theorem states that mixed and behavioural (and thus general) strategies have the same expressive power [Aum64].

These hypotheses have been inconspicuously challenged in recent papers. In this regard, the comparison between [BGG09] and [GS09] makes for an enlightening example. At first glance, these two papers look very similar: they both ponder the problem of the existence of almost-sure strategies in games where both players have (asymmetric) imperfect information. A closer examination reveals the differences: Bertrand, Genest, and Gimbert use general strategies, while Gripon and Serre use behavioural strategies; furthermore, in the latter paper, the players cannot observe their own actions. As a consequence, there are cases where the answer to the synthesis problem depends on which model is used. Consider for example the synchronous arena depicted in Figure 1, where Eve cannot distinguish vertices nor actions in the dashed area,  $\otimes$  is a losing sink state, and  $\odot$  is her “target”, for either a reachability or a Büchi condition.




Figure 1: Who wins?

With a behavioural strategy, Eve's strategy can only depend on the length of the play. At any even move, if her strategy is to play  $a$  with probability  $p$  and  $b$  with probability  $1 - p$ , Adam can answer by playing  $A$  with probability  $1 - p$  and  $B$  with probability  $p$ , so the odds of the token going to  $\alpha$  or to  $\beta$  are equal (they are worth  $p \cdot (1 - p)$  each). In the next step, no matter what  $\sigma$  advocates, the odds of the token going to  $\odot$  or to  $\otimes$  will again be equal. In the reachability game, this limits Eve's prospects to half chances. In the Büchi game, the probability that she wins drops to 0.

On the flip side, she has an almost-sure mixed strategy for both objectives: the natural “uniform” measure over the strategies of the form  $(aa|bb)^\omega$  guarantees that each sequence of two moves starting in the initial vertex has a probability of  $\frac{1}{2}$  to send the token to  $\odot$ , and a probability of  $\frac{1}{2}$  to send the token back to the initial vertex, no matter what Adam does. It cannot go to  $\otimes$ , as Eve never plays  $ab$  or  $ba$  from the initial vertex.

The arena of Figure 1 is synchronous, so any behavioural strategy can be emulated by a mixed one. If we remove this hypothesis, it is not always the case, as in the one-player game of Figure 2, where Eve is unaware of any action or vertex.




Figure 2: The D.U.I. game

As Eve observes nothing, her strategy is completely determined by what she does on the empty word  $\lambda$ . She has only two pure strategies:  $\lambda \rightarrow \text{stay}$  and  $\lambda \rightarrow \text{leave}$ . Both lead to  $\otimes$ , and so does any mixed strategy of the form  $\{p \cdot (\lambda \rightarrow \text{stay}), (1 - p) \cdot (\lambda \rightarrow \text{leave})\}$ . The behavioural strategy  $\lambda \rightarrow \{\frac{1}{2} \cdot \text{stay}, \frac{1}{2} \cdot \text{leave}\}$ , on the other hand, yields one chance out of four to reach  $\odot$ .

The case of games with perfect information and invisible actions is still open: there are mixed strategies which cannot be imitated by any behavioural one, so we cannot hope for a “generic” translation. But

that does not rule out the possibility of specific, objective-dependent constructions which would yield a different strategy with the same value.

## 4 Memory issues

A refinement of the synthesis problem asks that the controller uses only finite memory, as a natural requirement for implementability. Pure strategies with memory are defined in the following way:

**Definition 1** A pure strategy  $\sigma$  with memory  $M$  is a triple  $(\sigma^i, \sigma^u, \sigma^a)$  where  $\sigma^i \in M$  is the initial memory state;  $\sigma^u : (\Phi \times M) \rightarrow M$  is the memory update function, which maps a signal and a memory state to a new memory state and is called at each new observation of Eve; and  $\sigma^a : (\Phi \times M) \rightarrow \mathcal{X}$  is the next-action function, which maps a signal and a memory state to an action and is called at each step.

Notice that any pure strategy  $\sigma$  can be represented as a strategy with memory  $\Phi^*$ , with  $\sigma^i = \lambda$ ,  $\sigma^u = \cdot$  and  $\sigma^a = \sigma$ . A strategy has *finite memory* if  $M$  is a finite set, and is *memoryless* if  $M$  is a singleton.

Randomized strategies with (countable) memory are defined with similar tuples, except that some of their elements use randomization.

**Behavioural:** In a *behavioural strategy with memory  $M$* , the next-action  $\sigma^a : (\Phi \times M) \rightarrow \mathcal{D}(\mathcal{X})$  is randomized.

**Mixed:** In a *mixed strategy with memory  $M$* , the initial memory  $\sigma^i \in \mathcal{D}(M)$  is randomised.

**General:** In a *general strategy with memory  $M$* , the next-action, initial memory, and memory-update  $\sigma^u : (\Phi \times M) \rightarrow \mathcal{D}(M)$  are randomised.

The memory requirements can also depend of the type of strategy. In the game of Figure 1, for example, there is no almost-sure mixed strategy with finite memory (in the reachability game, there are  $\varepsilon$ -optimal strategies with finite (unbounded) memory; in the Büchi game, every mixed strategy with finite memory has value 0). However, the strategy we described can be realized by a general strategy with four memory states  $a_{\text{even}}$ ,  $a_{\text{odd}}$ ,  $b_{\text{even}}$  and  $b_{\text{odd}}$ : in the even memory state, she updates her memory at random to one of the odd states; in the odd states, she updates her memory to the corresponding even state; in all states, she plays the action corresponding to her memory state.

### 4.1 Concurrent safety games

In [dAHK98], de Alfaro, Henzinger, and Kupferman study the problem of concurrent reachability/safety games and establish the qualitative determinacy of these games, as well as several results on the nature (memory and randomization) of the strategies needed to achieve various objectives. In particular, they show that positive strategies for safety objectives require, in general, an infinite amount of memory. The proof is based on the famous “snowball game” of [KS81], which is pictured in Figure 3.

In this game, Adam loses if he never runs and Eve never throws, or if Eve happens to throw the snowball exactly at the moment he runs. It is clear that Adam has memoryless behavioural strategies with value arbitrarily close to one: if, at each step, he chooses to run with probability  $\varepsilon$ , he ensures a probability of winning of  $1 - \varepsilon$  (Eve’s best chance is to throw the ball right away). It is also clear that he cannot win almost-surely: if he has a positive probability of never running, Eve can keep the snowball forever; and if he has a positive probability of running at any step, Eve can thwart him by throwing the ball with probability  $\frac{1}{2}$  at each step.




Figure 3: The snowball game

By the *qualitative determinacy* of concurrent regular games, Eve has a positive strategy, *i.e.* a unique strategy which prevents Adam from winning almost-surely with any strategy. De Alfaro, Henzinger, and Kupferman use behavioural strategies, and argue that Eve needs infinite memory: the sequence  $(\sigma(\circ^i)(\text{throw}))_{i \in \mathbb{N}}$  must go to 0 but never reach it. It is clear that there are no positive mixed strategies with finite memory, as pure strategies with  $n$  memory states can only throw the snowball in the first  $n$  steps.

On the other hand, there is a general strategy with only 2 memory states: in the memory state *Never*, Eve keeps the snowball with probability 1, in the state *Eventually*, she throws it with probability  $\frac{1}{2}$ ; the memory never changes, and the initial memory state is chosen at random. This strategy prevents Adam from winning almost-surely, since he can never be sure that Eve is not in the memory state *Eventually*. In fact, this is the case in every finite concurrent safety game:

**Theorem 2** *In every finite concurrent safety game, Eve has a positive general strategy with memory 2 from her positive region.*

*Sketch of proof.* It follows from the analysis of the fix-points in [dAHK98] that there is a total preorder  $\prec$  on the vertices such that:

- the minimal vertices belong to the almost-sure region of Adam;
- for each non-minimal vertex  $q \in \mathcal{V}$ , there is an action  $\text{safe}_q \in \mathcal{X}$  of Eve such that, for any action  $y$  of Adam:
  - either for any vertex  $r \in \mathcal{V}$ ,  $\delta(q, \text{safe}_q, y)(r) > 0 \Rightarrow q \prec r$ ,
  - or there is an action  $x \in \mathcal{X}$  of Eve and a vertex  $r \in \mathcal{V}$  such that  $\delta(q, x, y)(r) > 0 \wedge q \not\prec r$

Notice that always playing the *safe* action is a pure and positional sure strategy for Eve in the maximal vertices (unless they are also minimal). For the vertices in between, we claim that the following strategy with two memory states is positive for Eve:

- the memory states are called *Sound* and *Chance*;
- each time the token goes to a new  $\prec$ -class, the memory state is updated to either *Sound* or *Chance* with equal probabilities; otherwise, the memory does not change;
- in the *Sound* state, Eve always plays the *safe* action of the current vertex;
- in the *Chance* state, Eve plays any action in  $\mathcal{X}$  with equal probabilities.

The situation is roughly the same as in the snowball game: if Adam's actions have no chance to go to a lower vertex against the *Sound* strategy, he will lose with probability  $\frac{1}{2}$ ; if he takes a risk at any point, there is a positive probability that Eve was in the *Chance* memory state all along, so he could end up in a greater vertex.  $\square$

In addition to the finite memory, the strategy described in the proof of Theorem 2 is simple, generic, and uses only uniform probabilities. By comparison, the description of a positive behavioural strategy is in general very complex and uses probabilities of unbounded precision.

## 4.2 Memory bounds for Muller games

Even in the elementary case of simple Muller games, it is not clear that the memory needs are the same for behavioural and general strategies. Recall that a simple arena is an arena with turn-based moves and perfect information for both players, and a Muller condition is a condition depending only on the set of colours visited infinitely often:

**Definition 3** A Muller condition on a set of colours  $\Gamma$  is specified by a subset  $\mathcal{F}$  of  $\mathcal{P}(\Gamma)$ . A play  $\pi$  satisfies  $\text{Muller}(\mathcal{F})$  if and only if the set of colours occurring infinitely often in  $\gamma(\pi)$  belongs to  $\mathcal{F}$ .

In such games, both players have pure optimal strategies with finite memory [BL69]. A follow-up problem is to determine, for a given Muller condition  $\mathcal{F}$  on a set of colours  $\Gamma$ , the necessary and sufficient amount of memory needed to define optimal pure strategies in *any* arena coloured by  $\Gamma$ . Gurevich and Harrington used the *latest appearance record* (LAR) structure of McNaughton to give a first upper bound of  $|\Gamma|!$  [GH82]. Zielonka refined the LAR into a tree, whose leaves could be used as memory [Zie98]. Finally, Dziembowski, Jurdzinski, and Walukiewicz showed that each player needs only as much memory as the number of leaves in some particular sub-trees, establishing tight and asymmetrical bounds for pure strategies [DJW97].

It is clear from their proof that mixing strategies does not help, since the other player can efficiently adapt their strategy in the witness arenas. This is not the case for behavioural strategies: Chatterjee, de Alfaro, and Henzinger observed that upward-closed winning conditions admitted memoryless strategies [CdAH04], leading to smaller upper bounds for arbitrary Muller conditions [Cha07]. Horn established even smaller tight bounds for general strategies [Hor09] (see Figure 4 for a graphical representation of the three bounds on a Zielonka (sub-)tree). However, Horn's upper-bound has not yet been proven (or refuted) for behavioural strategies.

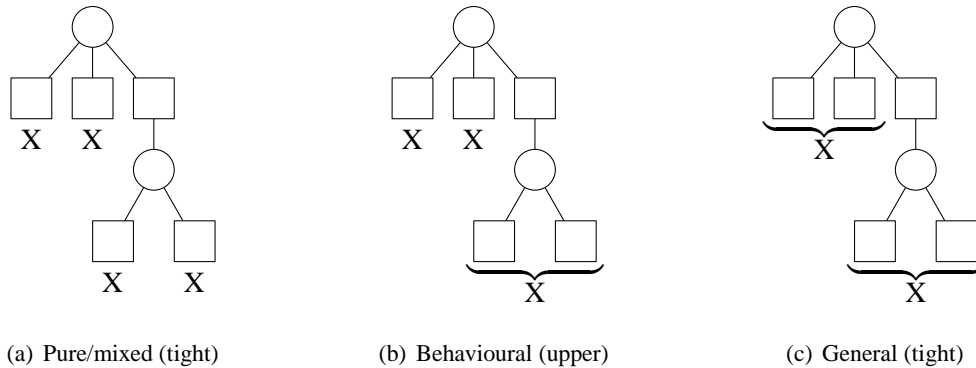


Figure 4: Memory bounds for simple Muller games



## 5 Discussion

We have compared three models of randomized strategies —behavioural, mixed, and general— in the context of graph games. Depending on the sub-case, we were able to expose variations in the amount of memory needed, the existence of finite-memory strategies, or even the values. In concurrent games with unobservable actions, the equivalence between the three models is still an open question.

In verification, the behavioural model has received most of the attention. Nevertheless, there is *a priori* nothing wrong with the other types of controllers. Furthermore, in several cases, general strategies can be much simpler than behavioural or mixed ones. On the other hand, general strategies are much less amenable to further analysis, as they introduce imperfect information. Even in simple safety games, one cannot compute the value of a general strategy —or even decide if it has positive value [GO09].

Each model has strengths and weaknesses, and we do not favour one over the others. Our point is rather to stress the importance of this initial choice, and to note that many memory-related problems which have been solved for behavioural strategies are still open in the mixed and general frameworks.

## References

- [Aum64] Robert J. Aumann. Mixed and Behavior Strategies in Infinite Extensive Games. In *Advances in Game Theory*, volume 52 of *Annals of Mathematical Studies*, pages 627–650. Princeton University Press, 1964.
- [BB08] Luciano Bello and Maximiliano Bertacchini. Predictable RNG in the vulnerable Debian OpenSSL package, the What and the How. In the *2nd DEF CON Hacking Conference*, 2008.
- [BGG09] Nathalie Bertrand, Blaise Genest, and Hugo Gimbert. Qualitative Determinacy and Decidability of Stochastic Games with Signals. In *Proceedings of the 24th Annual IEEE Symposium on Logic in Computer Science, LICS'09*, pages 319–328. IEEE Computer Society, 2009.
- [BL69] J. Richard Büchi and Lawrence H. Landweber. Solving Sequential Conditions by Finite-State Strategies. *Transactions of the American Mathematical Society*, 138:295–311, 1969.
- [CdAH04] Krishnendu Chatterjee, Luca de Alfaro, and Thomas A. Henzinger. Trading Memory for Randomness. In *Proceedings of the 1st International Conference on Quantitative Evaluation of Systems, QEST'04*, pages 206–217. IEEE Computer Society, 2004.
- [Cha07] Krishnendu Chatterjee. Optimal Strategy Synthesis in Stochastic Muller Games. In *Proceedings of the 10th International Conference on the Foundations of Software Science and Computational Structures, FoSSaCS'07*, volume 4423 of *Lecture Notes in Computer Science*, pages 138–152. Springer-Verlag, 2007.
- [dAH00] Luca de Alfaro and Thomas A. Henzinger. Concurrent  $\omega$ -regular Games. In *Proceedings of the 16th Annual IEEE Symposium on Logic in Computer Science, LICS'00*, pages 141–154. IEEE Computer Society, 2000.
- [dAHK98] Luca de Alfaro, Thomas A. Henzinger, and Orna Kupferman. Concurrent Reachability Games. In *Proceedings of the 39th Annual Symposium on Foundations of Computer Science, FoCS'98*, pages 564–575. IEEE Computer Society, 1998.
- [DHR08] Laurent Doyen, Thomas A. Henzinger, and Jean-François Raskin. Equivalence of Labeled Markov Chains. *International Journal of Foundations of Computer Science*, 19(3):549–563, 2008.
- [Dic55] Philip K. Dick. *Solar Lottery*. Ace Books, 1955.
- [DJW97] Stefan Dziembowski, Marcin Jurdziński, and Igor Walukiewicz. How Much Memory is Needed to Win Infinite Games? In *Proceedings of the 12th Annual IEEE Symposium on Logic in Computer Science, LICS'97*, pages 99–110. IEEE Computer Society, 1997.

- [GH82] Yuri Gurevich and Leo Harrington. Trees, Automata, and Games. In *Proceedings of the 14th Annual ACM Symposium on Theory of Computing, STOC'82*, pages 60–65. ACM Press, 1982.
- [GO09] Hugo Gimbert and Youssouf Oualhadj. Automates Probabilistes : Problèmes Décidables et Indécidables. Technical Report hal-00422888, Laboratoire Bordelais de Recherche en Informatique, CNRS UMR 5800, 2009.
- [GS09] Vincent Gripon and Olivier Serre. Qualitative Concurrent Games with Imperfect Information. In *Proceedings of the 36th International Colloquium on Automata, Languages and Programming, ICALP'09*, volume 5556 of *Lecture Notes in Computer Science*, pages 200–211. Springer-Verlag, 2009.
- [Hor09] Florian Horn. Random Fruits on the Zielonka Tree. In *Proceedings of the 26th International Symposium on Theoretical Aspects of Computer Science, STACS'09*, volume 3 of *Leibniz International Proceedings in Informatics*, pages 541–552. Schloß Dagstuhl, 2009.
- [KS81] Panganamala R. Kumar and Tzong-Huei Shiau. Existence of Value and Randomized Strategies in Zero-sum Discrete-time Stochastic Dynamic Games. *SIAM Journal on Control and Optimization*, 19(5):617–634, 1981.
- [Mar98] Donald A. Martin. The Determinacy of Blackwell Games. *Journal of Symbolic Logic*, 63(4):1565–1581, 1998.
- [MP92] Zohar Manna and Amir Pnueli. *The Temporal Logic of Reactive and Concurrent Systems: Specification*. Springer-Verlag, 1992.
- [MS98] Ashok P. Maitra and William D. Sudderth. Finitely additive stochastic games with Borel measurable payoffs. *International Journal of Game Theory*, 27(2):257–267, 1998.
- [PR89] Amir Pnueli and Roni Rosner. On the Synthesis of a Reactive Module. In *Proceedings of the 16th Annual ACM Symposium on Principles of Programming Languages, POPL'89*, pages 179–190, 1989.
- [TNJ07] Susan Thomson, Thomas Narten, and Tatuja Jinmei. IPv6 Stateless Address Autoconfiguration. RFC 4862 (Draft Standard), 2007.
- [Wil91] David Williams. *Probability with Martingales*. Cambridge University Press, 1991.
- [YL06] Tatu Ylonen and Chris Lonvick. The Secure Shell (SSH) Transport Layer Protocol. RFC 4253 (Proposed Standard), 2006.
- [Zie98] Wiesław Zielonka. Infinite Games on Finitely Coloured Graphs with Applications to Automata on Infinite Trees. *Theoretical Computer Science*, 200(1–2):135–183, 1998.